The *TV-Anytime* Forum

Specification Series: S-4

On:

Content Referencing

(Normative)

Document: SP004V11
Date: 14 April 2001

Series Overview

This is the fourth in a series of five “S-series” documents produced by the *TV-Anytime* Forum. These documents establish the fundamental specifications for the services, systems and devices that will conform to the *TV-Anytime* standard, to a level of detail which is implementable for compliant products and services.

As is common practice in such standardization efforts, these specification documents were preceded by requirements documents (“R-series”), which define the requirements for the *TV-Anytime* services, systems, and devices.

Congruent with the structure defined in *TV-Anytime*’s Call for Contributions (TV014r3), these specifications are parsed into three major areas, each described in a separate document of the series: Metadata (S-3), Content Referencing (S-4) and Rights Management (S-5). See the Call for Contributions for more detail on the derivation and background of these categories and their respective roles in the *TV-Anytime* standardization process.

The other two documents in the S-series are intended to define the environment and system architecture in which the standards in S-3, S-4, and S-5 are to be implemented. The first document in the series (S-1) provides benchmark business models against which the *TV-Anytime* system architecture is evaluated to ensure that the *TV-Anytime* standard enables key business applications. The next document in the series (S-2) presents the *TV-Anytime* System Architecture. These two documents are placed ahead of the other three for their obvious introductory value; S-1 and S-2 are both informative.

Although each of the S-series documents is intended to stand alone, a complete and coherent sense of the *TV-Anytime* system standard can be gathered by reading all five of the specification documents in numerical order.

Content Referencing Specification S-4
Document Revision History

Version	Status	Date	Author	Comments
WD 126	Working draft	30 May to 1 June 2000	Content Referencing Working Group	Drafted at the 6th TVAF meeting in New York City
TV050, r0	First draft of specification	26 - 28 July 2000	Content Referencing Working Group	WD126 approved by final plenary as TV050r0 at 7th TVAF meeting in Geneva, Switzerland
TV050, r1	First revision of specifications document	29 July – 9 August 2000	Alex Ashley - Content Referencing Group Editor (Philips)	Final editing by Content Referencing working group
TV050, r2	Final draft of first revision of specifications document	9 - 11 August 2000	Kellie McKeown (consultant)	Cosmetic editing and formatting
TV050, r3	Changes to TV050r2 in response to comments.	26 - 28 September 2000	Alex Ashley - Content Referencing Group Editor (Philips)	Alterations to specification by Content Referencing Working Group in response to comments to TV050r2.
SP004	Version 1 Content Referencing Specification.	2 - 13 October 2000	Nigel Earnshaw- Content Referencing Group Editor (BBC)	Cosmetic editing and formatting.
TV050r4	Provisional version of V1.1	16 February, 2001	Nigel Earnshaw	Provisional release V1.1
TV079	Changes to TV050r4 in response to comments.	27-30 March, 2001	Nigel Earnshaw- Content Referencing Group Editor (BBC)	Alterations to specification by Content Referencing Working Group in response to comments to TV050r4
SP004V1.1	Final Version 1.1 Content Referencing Specification.	2 - 14 April 2001	Nigel Earnshaw- Content Referencing Group Editor (BBC)	Cosmetic editing and formatting.

Table of Contents

Figures

Tables

About the TV-Anytime Forum

The global *TV-Anytime* Forum is an association of organizations which seeks to develop specifications to enable audio-visual and other services based on mass-market high volume digital storage in consumer platforms – simply referred to as *local storage*.

The *TV-Anytime* Forum was formed at an inaugural meeting held in Newport Beach, California, USA, on 27-29 September 1999. It has started work to develop open specifications designed to allow Consumer Electronics Manufacturers, Content Creators, Telcos, Broadcasters and Service Providers to exploit local storage.

As part of its formation, the *TV-Anytime* Forum has established four fundamental objectives for the organization, which are:

- The *TV-Anytime* Forum will define specifications that will enable applications to exploit local persistent storage in consumer electronics platforms.
- The *TV-Anytime* Forum is network independent with regard to the means for content delivery to consumer electronics equipment, including various delivery mechanisms (e.g. ATSC, DVB, DBS and others) and the Internet and enhanced TV.
- The *TV-Anytime* Forum will develop specifications for inter-operable and integrated systems, from content creators/providers, through service providers, to the consumers.
- The *TV-Anytime* Forum will specify the necessary security structures to protect the interests of all parties involved.

Member organizations from Europe, the USA, and Asia, are drawn from a wide variety of industries: Traditional Broadcasters, Internet Broadcasters, Content Owners, Service Providers, Telcos, Consumer Electronics Manufacturers, IT Industries, Professional Equipment Manufacturers, Component Manufacturers and Software Vendors.

The *TV-Anytime* Forum invites *participation* from all interested organizations. Membership is open to all who sign the Memorandum of Understanding (see Appendix A) and attend meetings. Meetings are held approximately every two months in Europe, the USA, and Asia.

For more information or to get involved with the work of the *TV-Anytime* Forum, visit the *TV-Anytime* Forum (www.tv-anytime.org) or contact:

Chairperson:	Simon Parnall (sparnall@ndsuk.com)
Vice-chairpersons:	Henry D. Chadwick (hdchad@us.ibm.com)
	Sakae Okubo (okubo@giti.or.jp)
Administrator:	Lynne Apodaca (lapodaca@inventures.com)
	(Tel: +1.925.275.6653)

1. Scope

The scope of this specification comprises the location independent identification of content, the location information, and the process of finding the location when an identifier is given. The content referencing process begins after a content item has been selected by a consumer through to, but not including, the actual acquisition of the desired content item.

Version 1.0 of the content referencing specification provides the mechanisms for location resolution in uni-directional network environments as defined in Business Model 1 of the *TV-Anytime* R-1 document. This version of the content referencing specification adds support for location resolution in a bi-directional environment to fulfill Business Models 2 and 3.

2. Glossary of Terms

Acquisition	Retrieval of content
Authority	Organization that creates CRIDs
Capture	Storing the acquired content (e.g., to local storage)
Content	Anything the viewer would like to access (movies, games, TV programs, radio programs, etc.)
Content Creator	Producers of content
Content Provider	Entity that acts as the agent for and is the prime exploiter of the content
Content reference	Pointer to a specific content item
Location Resolution	Process of establishing the address (location and time) of a specific content instance from its CRID
Locator	Time and place where a content item can be acquired
Metadata	Generally, data about content, such as the title, genre and summary of a television program. In the context of *TV-Anytime*, metadata also includes consumer profile and history data.
Resolution Handler	Functional unit that provides location resolution on a specific transport mechanism
Resolving Authority	Body which provides location resolution
Resolving Authority Record	Information needed for retrieving the location resolution data for the given authority
Service Provider	Aggregator and supplier of content which may include gateway and management roles

Abbreviations

CRID	Content Reference Identifier: identifier for content that is independent of its location
DVB	Digital Video Broadcasting: set of standards used for European digital TV broadcasting
DNS	Domain Naming System: system used on the Internet to register names that can then be mapped into IP addresses using a DNS server
EPG	Electronic Program Guide: means of presenting available content to the consumer, allowing selection of desired content
IP	Internet Protocol: generic name for the network protocols used on the Internet
IPR	Intellectual Property Rights
PDC	Program Delivery Control: European system used by some broadcasters to accurately indicate when programs start and stop
PDR	Personal Digital Recorder
SI	System Information: collection of information tables used in DVB
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
UTC	Coordinated Universal Time

3. Introduction

The purpose of content referencing is to allow acquisition of a specific instance of a specific item of content. For example, if a consumer sees an announcement on TV saying “there’ll be a new series of the Agatha Christie’s Murder Mysteries next year”, he/she may want to instruct their Personal Digital Recorder (PDR) to record the whole series, but cannot since he/she does not know when the episodes are going to be broadcast. In fact, the broadcaster may not know yet either. Still the viewer will want to make sure at this point that he/she does not miss the opportunity to acquire the content.

To provide the capability desired by the consumer, the ability is needed to refer to content (in this example a series of programs) independent of its location, whether that location is on a particular broadcast channel on some date and time, or on a file server connected to Internet, or wherever.

In this example, the PDR system would be provided with a reference for the series. In due time, the body who assigned the reference would provide the information required to link this reference to the individual episodes, and subsequently to a specific date and time for each episode so that the PDR would be able to acquire all of them.

This example demonstrates the purpose of content referencing – to provide the ability to refer to content independent of its location, and the ability to subsequently resolve such a reference into one or more locations where the content can be obtained.

Of course ‘*content’* can refer to many types of information. In addition to the television programs in the example above, it may include radio programs, audio tracks, MPEG-4 objects, scenes, images, music, etc.

4. Key Concepts and Features

The content referencing process begins after a content item has been selected by a consumer through to, but not including, the actual acquisition of the desired content item. In Figure 4-1 the scope of the content referencing process is shown. Some parts outside the scope of content referencing are included in Figure 4-1 to give it a global context.

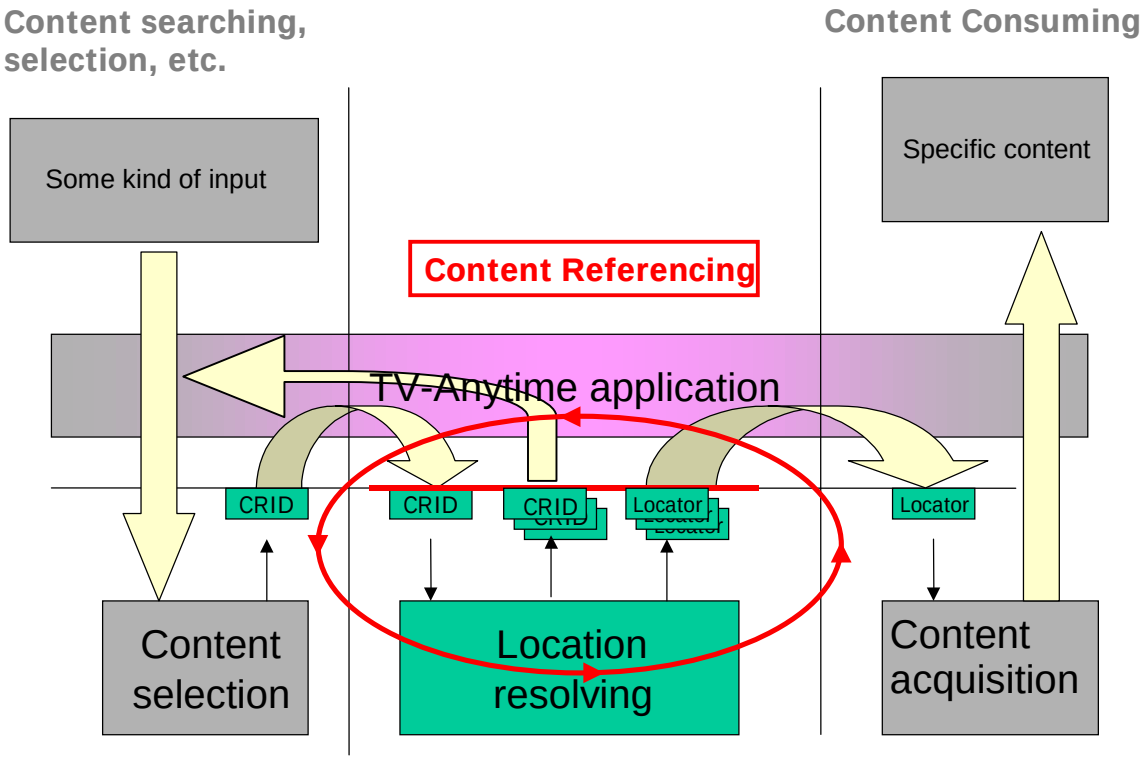

Figure 4-1: Content Referencing Environment

In this model we can identify three areas of interest:

1. The content selection process resulting in a Content Reference Identifier (CRID).
2. The location resolving service – the core of content referencing – resulting in one or more pointers to instances of that content (locators) or one or more content reference identifiers.
3. The retrieval process using one of these locators to subsequently acquire the content.

Before the process of content referencing can be employed, a selection process must supply a CRID. This CRID identifies a specific content item, but does not specify where that content item can be found. More detailed characteristics of the content, such as the price, the quality, the presence of commercials, etc., may not be known before the CRID is produced.

Given a CRID, there can still be numerous instances of the same desired content item (for example a broadcast may be repeated, a web site may be mirrored, the content item may be available from different content providers, via different networks, etc.) Some iteration may be required to narrow down the selection.

A content item that a CRID refers to might be a grouping of other content items, for example a CRID for referring to an entire series of programs.

The key concept of content referencing is the separation of the reference to a content item – the CRID – and information that is related to its retrieval.

The separation provided by the CRID enables a one-to-many mapping between content references and the locations of the deliverables.

Content is a general term. It is the context of a *TV-Anytime* service that determines what content will be delivered upon following a locator. In that sense a content instance is any deliverable that can be acquired (via a respective locator) by a consumer. New types of content delivery merely require new types of locators.

In the same way, it is the *TV-Anytime* location resolution service provider who declares what can be considered content items (e.g. programs, serials, etc.)

Content referencing, in conjunction with applications and additional metadata, may be able to deliver “first time” success, but may also involve some iteration enabling the overall system to deliver features including:

- Selection between alternatives
- Selection of sub-elements
- Selection between near matches
- Selection of time of delivery
- Selection ahead of time of release
- Selection based on coding quality
- Selection based on cost of delivery chain
- Selection based on acceded rights
- A point of reference – not necessarily unique – for the content item and any associated metadata.

Content referencing is:

- a form of content identification data and a specification for a set of rules for resolving this data into content locations from where this content can be retrieved.

Content referencing is not:

- the selection or retrieval process (it is in between);
- the definition of metadata used in any inquiry.

5. Location Resolution

Location resolution is the process of translating a CRID into other CRIDs or locators. Location resolution involves mapping a location-independent content reference (the CRID) to its location in time (e.g. scheduled transmission time in a broadcast system) and space (e.g. TV channel, IP address). In the context of this specification, locations in time and space are referred to as “locators.”

The process of location resolution may happen inside the PDR device (for example in a broadcast only system) or using a physically remote server (e.g., a server on the Internet.)

6. The Authority

An authority, as defined in this specification, is the body that creates a CRID, which they guarantee to be unambiguous. An authority will also provide the ability for the CRID to be resolved into locators or other CRIDs.

In most PDR implementations, there will be multiple authorities that the PDR must be able to distinguish between. To distinguish between authorities, each authority has a unique name. This *TV-Anytime* specification uses the Domain Naming System (DNS) to provide unique names for each authority.

The syntax of an authority name is:

<DNS name><name_extension>

<DNS name> is a registered Internet domain name. (See RFC1591 for DNS name registration[2].) The <DNS name> is case insensitive and must be a fully qualified name according to the rules given by RFC 1591.

<name_extension> is an optional string (beginning with a ';' character) to enable multiple authorities to use the same DNS name. All <name_extension> elements which share the same <DNS name> must be unique. The <name_extension> section is case insensitive.

Some example authority names are:

- www.broadcaster.com
- ISP.net
- www.commerce.com;electronics

7. The CRID

A CRID is the output of the search and selection process and is an unambiguous identifier that refers to a piece of content, however multiple CRIDs may refer to that same piece of content.

It is permissible for a CRID to resolve into one or more other CRIDs. This CRID to CRID functionality can be used for two purposes:

- A CRID can resolve into multiple CRIDs for grouping content items such as an entire series of programs.
- A CRID may also resolve into one or more CRIDs to allow one authority to refer to CRIDs of another authority.

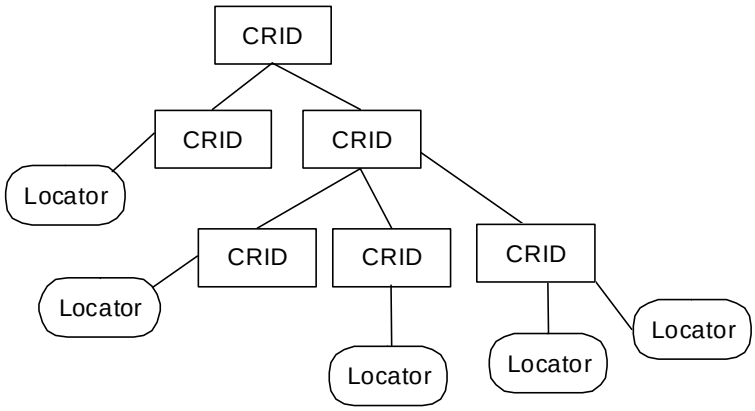

Figure 7-1: Example of a Tree-Structured CRID

The syntax of the CRID is:

CRID://<authority>/<data>

<authority> Uses the *TV-Anytime* authority naming rules given in section 6 to assure uniqueness.

<data> is a free format string that is Uniform Resource Identifier (URI) compliant, and is meaningful to the authority given by the <authority> field. The <data> portion of the CRID is case insensitive.

In its entirety, the CRID is URI compliant. (See RFC2396 for URI compliance specification [1].) As per RFC2396, the CRID:// part of the syntax is case insensitive.

Examples of syntactically valid CRIDs are:

CRID	Description
CRID://company.com/foobar	CRID created by “company.com” authority, with a data part of “foobar”
crid://broadcaster.co.jp;comedy/wibble	CRID created by “broadcaster.co.jp;comedy” authority, with a data part of “wibble”

Figure 7-2: Examples of CRID

8. Locators

An instance of content may be located on various media such as local storage, live broadcast stream, data broadcast stream, data file on the Internet and data stream via the Internet.

A locator specifies a location, and possibly time of availability, where a content item can be acquired. There will be many formats of locators as there are many different means by which a PDR can acquire content. It is a requirement of a locator to ensure that it is possible for the PDR to parse enough of the locator to be able to decide if it has the ability to use the relevant transport mechanism.

The locator will be parsed and used by media-dependent methods to identify the content location and to acquire the content using the media or transport specific protocol. For example, a DVB locator will contain location parameters for a DVB stream, such as transport stream ID, service ID, table ID and event ID.

The syntax for a locator is:

<transport mechanism>:<transport system specific>

<transport mechanism> shall be unique for each mechanism. The string “CRID” shall not be used as the name for a <transport mechanism>.

<transport system specific> will be defined by the creator of the <transport mechanism>.

In its entirety, the locator is URI compliant. (See RFC2396 for URI compliance specification [1].)

For each <transport mechanism>, there will be only one format for the syntax of the <transport system specific> section.

The <transport system specific> section shall provide the following information:

- Location – This provides the location where the content can be acquired. It is possible that a PDR can receive content from many different providers which all share the same <transport mechanism>. For this reason, it is a requirement of a TV Anytime locator that it is unambiguous between multiple providers using the same <transport mechanism>.
- Type of availability – It is possible that some schemes will be used for both schedule-based and on-demand acquired content. Content that is available at a specific time at a specific location (e.g., a broadcast TV program, a web–cast) is schedule based. Schedule-based content must be retrieved at the time given by the locator. Content that can be received at any time between two limits (e.g. content that is on a server for one month) is on-demand based. On-demand based content can be acquired at any time that it is available.

For schedule-based content:

- Start time – This provides information about when the content is scheduled to start. It is required that start time be unambiguous with respect to local time zone as a PDR may be able to receive content from many different time zones.
- Duration of content – The length in time of the content.

For on-demand content:

- Start of availability – This optional field, when present gives the first moment in time when the content is available. This time shall be unambiguous in terms of time zone as a PDR may be able to receive content from multiple time zones.
- End of availability – This optional field gives the first moment in time that the content will become unavailable. This time shall be unambiguous in terms of time zone as a PDR may be able to receive content from multiple time zones.

In defining a syntax for the <transport system specific> section of a locator associated with a <transport mechanism>, there is an assumption about the environment the PDR exists within.

For each <transport mechanism> a PDR will need a certain amount of information in order to receive content from this system. This information may be provided within the transport mechanism, or by any other means appropriate to the target PDR.

As an example of the information within the environment, a PAL Western European locator might use the network and channel identifier within its syntax. The mapping of network and channel identifier to a physical channel uses information carried in the vertical blanking interval.

The transport mechanism may provide a more accurate timing system than the start time, which the PDR may decide to use for accurate content capture (e.g. PDC information, DVB event IDs.)

9. The Resolving Authority Record

The resolving authority record (RAR) is an essential element in the location resolution process for both the unidirectional and bi-directional networks. It provides the information needed for retrieving the location resolution data for a given authority.

Each resolution authority will require one or more resolving authority records to exist in the PDR for location resolution to take place. Each resolving authority record will need to be placed inside some sort of transport specific container, which allows the PDR to know that this is a resolving authority record.

In the case of multiple records for the same authority, for each location resolution to be done, the PDR can choose to use any of them.

9.1 Resolving Authority Record Structure

This specification gives two forms of Resolving Authority Record structure described in the tables below. The field semantics follow Table 9-2.

A version 1.0 *TV-Anytime* compliant authority record has the following structure:

Field Name	Field Format
Syntax Version	One byte unsigned integer containing major version number One byte unsigned integer containing minor version number The following table is a v1.0 table, which is encoded using major number equals one, minor number equals zero.
Resolution Provider	One byte containing length of resolution provider name. 'n' characters encoded using ANSI X3.4-1986 (R1997) [3]
Authority name	One byte containing length of authority name 'n' characters encoded using ANSI X3.4-1986 (R1997) [3]
Class	One byte which contains either 'P' or 'S' encoded using ANSI X3.4-1986 (R1997) [3]
Version number	Four byte unsigned integer, most significant byte first
URL	One byte containing length of URL 'n' characters encoded using ANSI X3.4-1986 (R1997) [3]
First valid date	Two byte unsigned integer for year, most significant byte first One byte unsigned integer month (1..12) One byte unsigned integer day (1..31) One byte unsigned integer hour (0..23) One byte unsigned integer minute (0..59)
Last valid date	Two byte unsigned integer for year, most significant byte first One byte unsigned integer month (1..12) One byte unsigned integer day (1..31) One byte unsigned integer hour (0..23) One byte unsigned integer minute (0..59)
Weighting	One byte 2s complement signed integer

Table 9-1: v1.0 Resolving Authority Record

A version 2.0 *TV-Anytime* compliant authority record is defined in Table 9-2 below.

Field Name	Field Format
Syntax Version	One byte unsigned integer containing major version number One byte unsigned integer containing minor version number The following table is a v2.0 table, which is encoded using major number equals two, minor number equals zero.
Version number	Four byte unsigned integer, most significant byte first
Class	One byte which contains either 'P' or 'S' encoded using ANSI X3.4-1986 (R1997) [3]
Weighting	One byte 2s complement signed integer
First valid date	Two byte unsigned integer for year, most significant byte first One byte unsigned integer month (1..12) One byte unsigned integer day (1..31) One byte unsigned integer hour (0..23) One byte unsigned integer minute (0..59)
Last valid date	Two byte unsigned integer for year, most significant byte first One byte unsigned integer month (1..12) One byte unsigned integer day (1..31) One byte unsigned integer hour (0..23) One byte unsigned integer minute (0..59)
Reserved	16 bytes
Resolution Provider	One byte containing length of resolution provider name. 'n' characters encoded using ANSI X3.4-1986 (R1997) [3]
Authority name	One byte containing length of authority name 'n' characters encoded using ANSI X3.4-1986 (R1997) [3]
URL	One byte containing length of URL 'n' characters encoded using ANSI X3.4-1986 (R1997) [3]
Reserved Length	Two byte unsigned integer, most significant byte first.

Table 9-2: v2.0 Resolving Authority Record

Syntax Version -- The version of the table format. Major numbers are changed when the contents of the RAR change in such a manner that a parser conforming to a previous major number would be unable to parse this new syntax. A change in the minor version number may add new fields to the RAR, but these fields will be added in such a manner that a parser conforming to the major number given in this field will still be able to parse this RAR.

Resolution Provider -- The name of the body that is providing location resolution. It is possible that different bodies are providing location resolution for a single authority, for example a broadcaster might be providing a secondary resolution service for a content creator. These different location resolution providers need to be able to identify themselves for such purposes as updating their resolving authority records. The name of the resolution provider follows the naming rules given in section 6.

Authority name -- The name of the *TV-Anytime* location resolution authority as described in Section 6.

Class -- The class field defines whether this authority record defines a resolution authority that can resolve all CRIDs for this authority name (class = primary) or only resolves some CRIDs for this authority name (class = secondary.) The reason for providing secondary

resolving authority records is to allow one resolution provider to provide a partial resolving service for a content referencing authority. An example of this is a broadcaster providing partial resolution of another authority (such as the content creator) for the content they broadcast. For a primary class resolving authority record, the class field will contain the character 'P' encoded using ANSI X3.4-1986 (R1997). For a secondary class resolving authority record, this field will contain the character 'S' encoded using ANSI X3.4-1986 (R1997).

Version number -- A number that is incremented each time the resolution provider wishes to update its resolving authority records for a given authority name. The set of authority records the PDR should update is based on the combination of the authority name and the resolution provider. When a new version number is received for an authority by a resolution provider, all old resolving authority records for this authority name and resolution provider combination will be discarded by the PDR. When the version number reaches 2^{32}-1, the next version number shall be zero. Tables are considered equivalent if they have the same values of Resolution Provider, Authority name, Version number and URL.

URL -- The URL field points to the location where resolution information can be found. The URL could point to a broadcast stream, or to a server on the Internet or any other place where location resolution information can be found. The syntax of the URL is that of the locator as given in section 8.

First valid date -- The first date when this authority can be used, using Universal Time Coordinates (UTC) as the time reference. The year field includes the century (e.g., to specify a date in the year 2001 this field should contain 2001.) Seconds are always assumed to be zero, and so not included in this field.

Last valid date -- The first date when this authority cannot be used, using UTC as the time reference. The year field includes the century (e.g., to specify a date in the year 2001 this field should contain 2001.) Seconds are always assumed to be zero, and so not included in this field.

The reason for providing start and end dates for resolution is so that resolution providers can move their resolution URLs and be sure all PDRs have switched to the new URL once the last valid date of the old resolution record has passed.

Weighting – The weighting field is used to give a hint to the PDR as to the order to try multiple records for an authority from the same resolution provider by providing the largest weighting number to the URL that should be tried first. The weighting field is only used to provide ordering between resolution provider records for the same combination of resolution provider and authority name and not for ordering one provider over another.

Reserved Length - Number of bytes following this field before the end of the RAR.

9.2 Resolving Authority Record Examples

An example of a v1.0 resolving authority record is given in Figure 9-1 below.

Field Name	Field contents	Description
Syntax Version	1 0	Syntax v1.0
Resolution Provider	15 "tva.resprov.com"	Length of tva.resprov.com Name of resolution provider
Authority name	10 "autnam.com"	Length of autnam.com Authority name
Class	'S'	Secondary class
Version number	0,0,3,232	Version number 1000
URL	32 "http://www.resprov.com/lr/autnam"	Length of URL URL
First valid date	7,208 9 26 9 30	Year (2000) Month (9) Day (26) Hour (09) Minute (30)
Last valid date	7,208 9 28 18 0	Year (2000) Month (9) Day (28) Hour (18) Minute (00)
Weighting	1	1

Figure 9-1: Example of v1.0 Resolving Authority Record

An example of a v2.0 resolving authority record is given in Figure 9-2 below:

Field Name	Field contents	Description
Syntax Version	2 0	Syntax v2.0
Version number	0,0,3,232	Version number 1000
Class	'S'	Secondary class
Weighting	1	1
First valid date	7,208 9 26 9 30	Year (2000) Month (9) Day (26) Hour (09) Minute (30)
Last valid date	7,208 9 28 18 0	Year (2000) Month (9) Day (28) Hour (18) Minute (00)
Reserved	0,0,0,0,0,0,0,0, 0,0,0,0,0,0,0,0	16 reserved bytes
Resolution Provider	15 "tva.resprov.com"	Length of tva.resprov.com Name of resolution provider
Authority name	10 "autnam.com"	Length of autnam.com Authority name
URL	32 "http://www.resprov.com/lr/autnam"	Length of URL URL
Reserved Length	0,0	Zero reserved bytes after this field

Figure 9-2: Example of v2.0 Resolving Authority Record

10. Protocols for Location Resolution

10.1 Common Features and Interfaces

This section defines the features common to location resolution using uni-directional or bi-directional networks. It is included here to provide a context for the details that follow in sections 10.2 and 10.3, and provides the descriptions of the terms that will be used in these sections.

10.1.1 Resolution Handler

In the *TV-Anytime* environment, content referencing services can be delivered through various delivery systems such as IP networks or broadcast TV. Figure 10-1 shows the concept of a modular resolution system, with multiple resolution handlers as required for specific location resolution transport mechanisms.

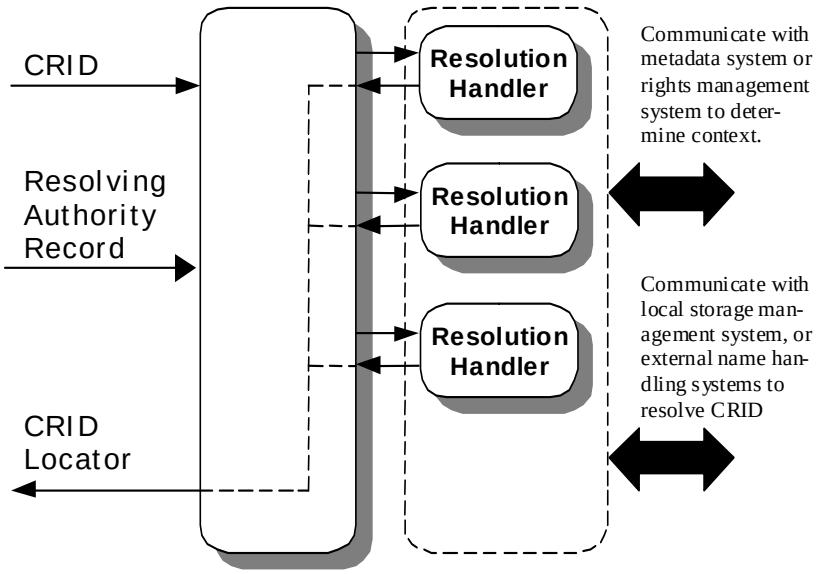

Figure 10-1: CRID Resolution Architecture

Figure 10-1 provides a network-transparent method for resolution, with multiple resolution handlers providing the network and protocol-dependent CRID resolution. For example, one resolution method is resolving CRIDs locally, by co-operating with the local storage management system to resolve the location of the locally stored content. Another resolution system might resolve a CRID using external name handling systems via a back channel or an Internet connection. Another system may refer to System Information (SI) tables, which contain mapping tables between CRIDs and locators and are transported in a digital broadcast stream.

It is envisaged that the resolution handler is an extensible resolution mechanism so that in the future when a new location resolution transportation method is developed, it is possible to extend the resolution process by implementing and plugging in a new resolution handler.

Steps in the content referencing process are now summarized.

1. The CRID is used to decide which resolution handlers to invoke to resolve this CRID.
2. The resolution request is forwarded to the appropriate resolution handlers.
3. Each chosen resolution handler tries to resolve the CRID into locators or another set of CRIDs. The resolution process depends on the implementation of resolution handler. As part of this resolution process, the resolution handler may need to communicate with an

external system. Some example processes between a resolution handler and an external system are:

i. Resolve the CRID using a mapping table located in the PDR. This method is appropriate for locally recorded contents or cached information from broadcast or IP networks.

ii. Resolve the CRID using the broadcast stream.

iii. Resolve the CRID via Internet or back channel.

10.2 Location Resolution Over Unidirectional Networks

This section of the specification defines the common features of resolution handlers that work with uni-directional networks. Each resolution handler using a uni-directional network will have a dynamic behavior similar to the one given in the flow chart Figure 10-2.

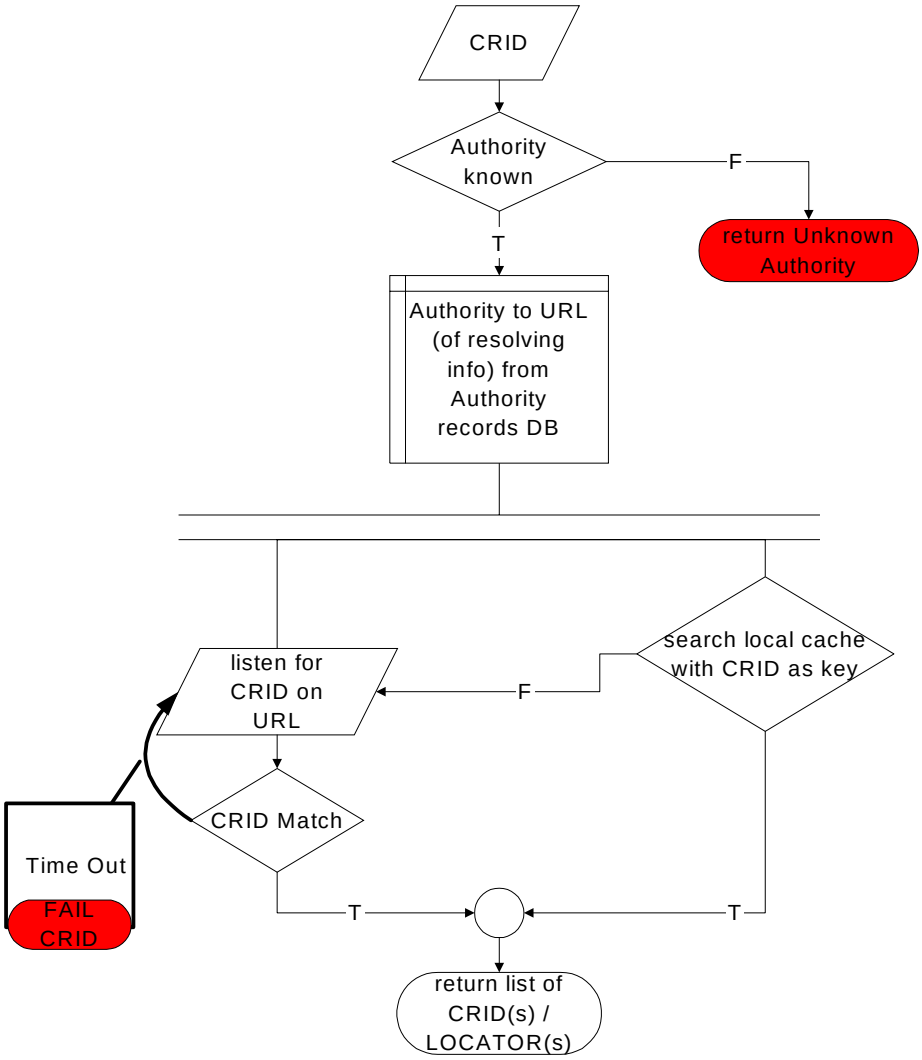

Figure 10-2: Dynamic Behavior of Resolution Handler using Unidirectional Network

The first step of location resolution in a unidirectional system is for the PDR to learn where it can receive location resolution information. This location is provided by the resolving authority record, which will need to be broadcast to the PDR at some known location. Failure to find any authority records for the given CRID's authority will result in a failure to resolve CRIDs created by that authority.

Once this resolving authority record has been located, the PDR will know where to listen for location information for the given CRIDs' authority (by using the URL field of the appropriate authority record.)

The PDR will have to choose one or more resolution handlers to use for resolving the desired CRID, as it will need to pick the resolution handlers that can understand the protocols used to carry the location resolution information. An example of this choice is the PDR choosing the DVB resolution handler if the resolution record says that resolution information is being sent on a DVB transport stream.

It is likely that a PDR implementation will also use the local resolution handler in case the content the CRID refers to is already available locally.

The information that will be sent in the unidirectional location resolution stream shall be in the form of a table that consists of CRID to message mappings. Each input CRID will output to a message that shall contain a status field. When the status field contains a value to indicate that the input CRID is valid, the message shall contain one or more CRIDs, or one or more locators.

A *TV-Anytime* compliant unidirectional location resolution stream shall contain a stream of matched pairs of:

CRID	Message

Each message shall at least contain:

Field	Description
Status	"CRID is resolved" (resolution list follows) "discard CRID" (E.g. CRID is no longer valid) "resolve after date <*xxx*>" (keep CRID, try later)

If Status = CRID is resolved:

Field	Description
Acquisition directive	"all" (all items of the following list must be acquired) "any" (any item from the following list may be acquired as they are alternative locations for the same content)
A list of CRIDs or **a list of Locator(s)**	CRIDs will conform to the syntax given in section 7 Locators will conform to the syntax given in section 8
Resolution complete	Is the list complete? yes (CRID is completely resolved) (e.g. This is the last episode of the series) no (CRID might resolve into more items at a later date)
Re-resolution date	Date after which the PDR should attempt to re-resolve the CRID. This field is only meaningful when the Resolution Complete flag is set to "no". This date shall be unambiguous with respect to timezone

If Status = 'Resolve after date'

Field	Description
Date	The date and time on or after which the PDR should try to re-resolve the CRID. This field shall be unambiguous with respect to timezone.

Table 10-1: Location Resolution Message Format for Unidirectional System

The following table describes the acquisition behavior of a PDR in response to the status flags

Acquisition Directive	Resolution Complete	Description
All	No	Acquire all items in list and await further items.
All	Yes	Acquire all content items, after which acquisition of this CRID is complete.
Any	No	Select any of current item list (after which the acquisition is complete) or await additions to list.
Any	Yes	Select any one of the list items after which the acquisition is complete

Table 10-2: Interpretation of Status Flags

10.2.1 Guidance on use of Resolution Status Flags

CRIDs that when resolved translate into one or more CRIDs, with the acquisition directive of "all" and have the resolution complete flag set to "no" can be used for grouping of content that changes over time, for example a TV series. Such a group CRID may continue with its resolution complete flag set to "no" for a long period of time if the series has no planned end. It is permissible for a PDR to allow the user to view the content the PDR has acquired for this incomplete group CRID.

A CRID that resolves into one or more locators should not be used for an ongoing group (such as a series) as the PDR will assume that when the acquisition directive is set to "all" it must acquire all the parts of the content specified by the list of locators before that content is fully acquired, and hence viewable by the user.

Resolution type	Acquisition directive	Description of PDR behavior
CRID to CRID(s)	All	All result CRIDs should be acquired. Each result CRID can be considered either as its own content item or as part of the group.
CRID to CRID(s)	Any	Any of the result CRIDs can be acquired as they are considered equivalent by the authority that created this CRID.
CRID to Locators(s)	All	All items must be acquired before the content is complete. It is an implementation issue as to whether a PDR will allow viewing of incomplete content.
CRID to Locators(s)	Any	Select any one of the locators as they are considered equivalent by the authority that created this CRID.

Table 10-3: PDR behaviour in response to acquisition directive

It is an implementer's option as to whether the PDR will always go to the unidirectional stream for location resolution, or provide some local caching mechanism. This caching mechanism might be to cache resolved CRIDs, or to cache the unidirectional stream in case it is needed later. This cached information might be used by the resolution handler that handles the unidirectional stream or by creating another resolution handler that uses locally cached data.

10.3 Location Resolution Over Bi-directional Networks

In order for a PDR to use the location resolution services over bi-directional networks it is necessary to define a protocol to allow the PDR to initiate a connection and then transfer requested data between itself and a resolution service located at a remote server.

This section specifies how a PDR can discover the location of such a server on a bi-directional network and the TVA protocol to achieve the appropriate TV-Anytime data transfers over such a network.

This section does not specify how content is retrieved over a bi-directional network.

10.3.1 Generic Bi-directional Resolution Server Discovery

This discovery stage may not be required by all network implementations.

Given a CRID to resolve, the first step is to find a server that might be able to resolve this CRID. The process of server discovery is based on using the authority name from the CRID.

This specification makes the assumption that the CRID the PDR has been asked to resolve is from an authority completely new to the PDR, so that it has no prior knowledge of where to get this CRID resolved. In an actual implementation there may be some sort of caching of previous server discoveries, but the exact nature of this caching is implementation specific.

In the following example, an intermediate server is contacted to discover the address of the location resolution server.

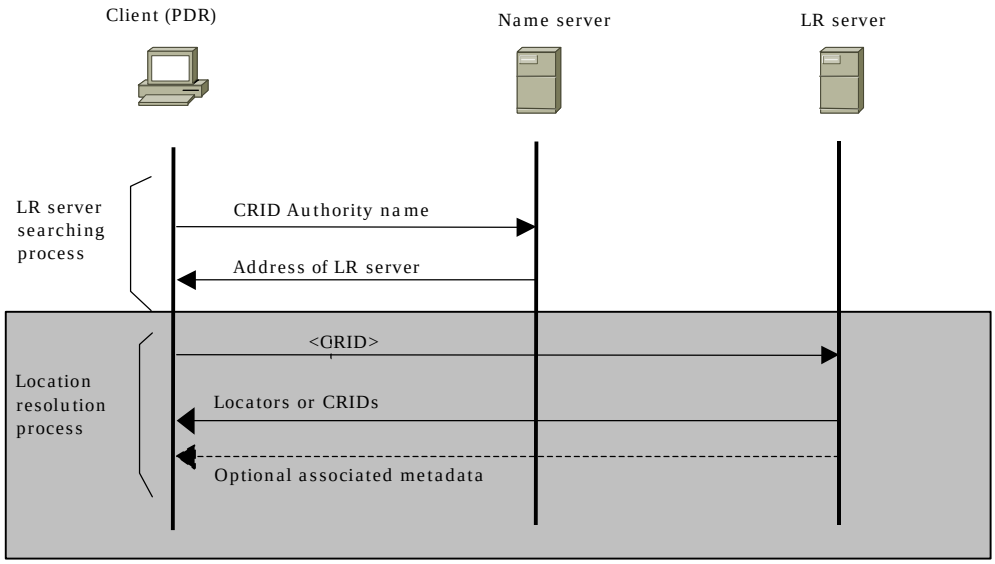

Figure 10-3: Example of Location Resolution Server Discovery

10.3.2 Generic Bi-directional Request to Resolution Server

Once the resolution server has been discovered, the next step is to communicate with this resolution server. The inputs to the server will be a list of one or more CRIDs to resolve, and optional flags to specify how the response should be created.

The TV Anytime defined optional flags are:

1. SubmittedCRID
2. Result

Value of SubmittedCRID flag	Description
0	No descriptive metadata about the CRIDs being submitted shall be returned with the resolution information.
1	Instances of the ProgramInformationTable or GroupInformationTable schemas shall be returned which describe the submitted CRIDs, if the location resolution server has this information.
All other values	Reserved

Table 10-4: Submitted CRID Flag Description

Value of Result flag	Description
0	No descriptive metadata about the results from the location resolution shall be returned with the resolution information.
1	If the submitted CRID resolves into further CRIDs, instances of the ProgramInformationTable or GroupInformationTable schemas shall be returned for each resolved CRID, if the location resolution server has this information. If the submitted CRID resolves into locators, instances of the ProgramLocationTable schema shall be returned for each locator, if the location resolution server has this information.
All other values	Reserved

Table 10-5: Result Flag Description

10.3.3 Generic Bi-directional Response From Resolution Server

The resolution server will respond with one of three possible types of information:

1. The result from resolving the CRID. The response will contain instance(s) of XML schemas defined in the TV Anytime Metadata specification SP003 [11] and in Appendix A of this specification. The permitted response schemas are listed later in this section.
2. A Resolving Authority Record (RAR). The PDR should store this RAR using the rules given in the uni-directional model, and then contact the server given by the URL field of the RAR
3. A re-direct. The resolution server will return a message that gives the address of another location resolution server to contact.

In case (1) and case(3) where a PDR does not receive a RAR, the PDR shall assume the location resolution server is a primary class server, and follow the appropriate rules given in this specification for primary class resolution servers.

The permitted instances of the XML schemas defined in SP003 [11] and in Appendix A of this specification are:

- GroupInformationTable
- ProgramInformationTable
- ProgramLocationTable
- ContentReferencingResult

- CRIDResult
- LocationsResult

When an XML instance document is returned by the location resolution server, instances of either ContentReferencingResult, CRIDResult or LocationsResult shall be returned for each CRID submitted. When indicated by the SubmittedCRID and Result flags, instances of GroupInformationTable, ProgramInformationTable and ProgramLocationTable may also be returned.

The results may be returned in any order the location resolution server desires and not necessarily the order in which the PDR specified in the request to the server.

10.3.4 Dynamic Behaviour of PDR and Location Resolution Server

For the case where the PDR is connected to a bi-directional network there is a need to specify some aspects of the dynamic behaviour of location resolution requests in a TV Anytime compliant PDR. This is required in order to avoid large numbers of PDRs all trying to contact the same location resolution server at the same moment in time.

As well as a PDR having a prescribed manner to avoid excess loading of a location resolution server, it is advisable that a location resolution provider takes certain steps in their implementation to aid the reduction of large instantaneous loads on their servers. When a server returns a response that indicates a CRID should be resolved again at a future date, it is advisable for the server not to return the same time and date information to all clients requesting this CRID, as they will then all try to contact the server at the same moment in time.

It should be noted that a PDR, in addition to receiving location information from a bi-directional network, might also be receiving information from one or more broadcast networks. When receiving information from a broadcast chain, there is the potential for many PDRs to all receive the same time and date for re-resolution.

In order to reduce the effect of instantaneous server loading, the PDR will need to implement a dynamic behaviour that does not cause overloading of a location resolution server by frequent repeated access to the same server.

10.3.4.1 Requirements for PDR dynamic behaviour

1. When the date and time for re-resolution arrives, the PDR shall wait a further random amount of time before contacting the location resolution server. This is required in order to reduce the chances of the location resolution server being swamped by many PDRs contacting it at the same moment in time.
2. If the location resolution server returns a re-resolution time and date that is in the past, the PDR shall wait a random amount of time before contacting the location resolution server again.
3. If the location resolution server is unavailable, the PDR shall assume the response for the CRIDs being resolved is “resolve again at a later date”. The date and time at which to re-resolve is based on the current date and time plus a random amount of time.
4. When a location resolution server returns information to specify that the CRID is unknown and this location resolution server is a secondary class server, and the PDR wishes to try this server again, the PDR shall wait a random amount of time before trying again.
5. When a location resolution server returns information to specify that the CRID is unknown and this location resolution server is a primary class server, the PDR should cease trying to resolve this CRID as it will never be resolvable.
6. The random time interval generator algorithm inside a given manufacturer’s PDR should not be configured in such a manner that all their PDRs have identical random time

interval sequences. Testing compliance to this requirement is not specified in this document.

7. The standard deviation of the random time interval generator shall be at least 10 minutes. Testing compliance to this requirement is not specified in this document.

Appendix B describes a candidate for the dynamic behaviour of the PDR when communicating with a remote location resolution server.

10.3.5 TCP/IP Based Resolution Server Discovery

As the <DNS name> part of the authority name is a name that is a registered Internet domain name, the mechanisms defined for DNS name lookup can be used as part of the server discovery phase.

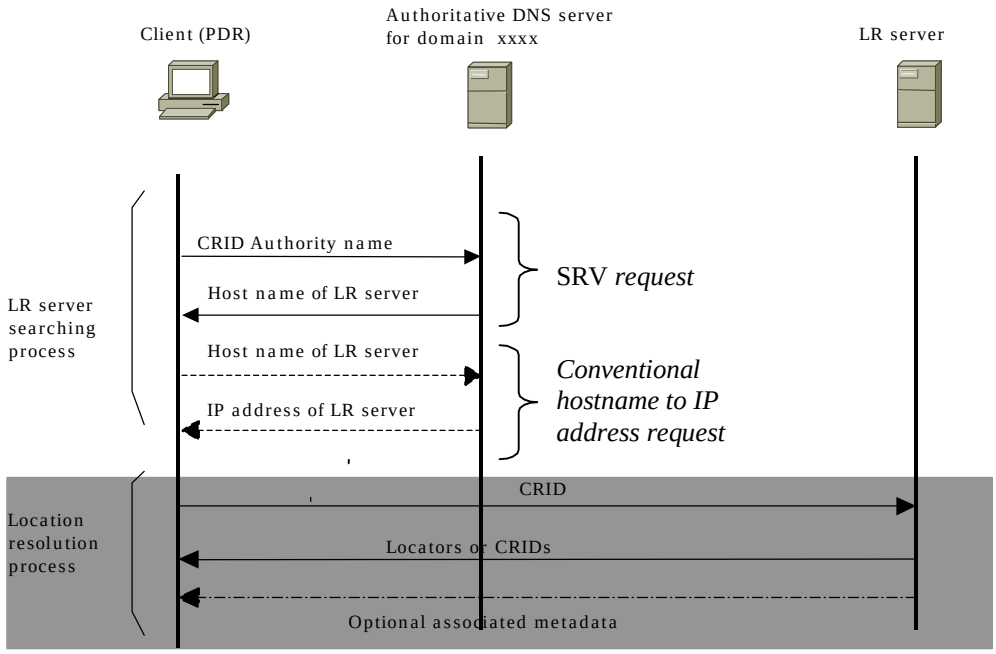

Figure 10-4: Stages of TCP/IP Based Resolution Server Discovery

10.3.5.1 Internet Service Discovery Query (RFC2782)

RFC2782 [8] provides an expansion of the DNS [2] system that is used to allow Internet connected machines to find mail servers. In addition to being able to search for mail exchange (MX) records, it is also possible to search for service (SRV) records.

An RFC2782 compliant query is made up of several parts, namely:

_Service._Protocol.Name

For example a query for an HTTP server for example.com would be "_http._tcp.example.com". The DNS server will respond with the host name and port number corresponding to the network location at which the requested service can be found. In the previous example, the return might be "webserver2.example.com on port 80".

10.3.5.2 TV-Anytime Location Resolution Service Query

The name for the TV Anytime location resolution service is "lres" which is a shortened version of "location resolution". The use of a shortened name is adopted, as there is a limit of 512 characters for a DNS response in some DNS client implementations.

The complete query name will look like:

_lres._tcp.*<name_extension segment from CRID authority>***.***<DNS segment from CRID authority>*

For example, given the CRID "crid://example.com;europe/9afc2", the query string would be "_lres._tcp.europe.example.com" which would be sent to a DNS server that provides lookups for "example.com".

Another example, this time without a name_extension segment in the authority name, would be: given the CRID "crid://example.co.uk/9afc2", the query string would be "_lres._tcp.example.co.uk" which would be sent to a DNS server that provides lookups for "example.co.uk".

10.3.6 TCP/IP Based Request to Resolution Server

The protocol for sending a request to the location resolution server is based on the HTTP protocol [4]. The format of the query string shall follow that generated by the submission of an HTML form using a GET request:

*http://<Path_to_server_script>***?***[key=value]***&***[key=value]*

where the key/value pair is repeated as required. The key is case sensitive and shall be represented using the exact case given for each key as specified in this section of the specification.

The precise specification for encoding the key-value pairs into the HTTP URL is given in section 17.13 of the HTML 4 specification [9]. It is the option of the location resolution service provider to implement this service using any server side technology they wish (CGI scripts, java servlets etc).

Each key shall be one of:

Key	Description	Allowable value
CRID	The CRID to resolve	A CRID inside quotes
SubmittedCRID	Used to specify whether metadata on the CRID being resolved is desired. See Section 10.3.2 for semantics of this flag.	A number. See Section 10.3.2 for allowed values.
Result	Used to specify whether metadata is required for each of the results from resolving this CRID. See Section 10.3.2 for semantics of this flag.	A number. See Section 10.3.2 for allowed values.

Table 10-6: Key definitions for HTTP URL encoding

It is permissible to resolve multiple CRIDs in one HTTP request by using multiple "CRID" keys in the URL, but the "SubmittedCRID" and "Result" keys may only be specified once in a request as they affect all CRIDs being resolved.

For the first connection to a location resolution server after the DNS based server location phase, the <path to server> is the hostname, a colon character, the textual representation of the port number followed by a slash. If the port number is 80, the colon and port number may be omitted.

For example, if the DNS server returned host name "computer2.example.com" on port 1234, the <path to server> would be

computer2.example.com:1234/

When a location resolution server provides a re-direct using an HTTP redirect, the <path to server> is the URL returned by the "Location" header of the HTTP redirect response.

For example, if the HTTP response was

Location http://redirect.example.com/tva/lr

The <path to server> would be

redirect.example.com/tva/lr

When a location resolution server provides a re-direct using the RAR, the <path to server> is the URL field from the RAR.

For example, if the URL field of the RAR contained the value

http://kaas.example.nl/scripts/resolution.cgi

the <path to server> would be

kaas.example.com/scripts/resolution.cgi

Examples of valid complete URLs are:

http://computer2.example.com:1234/?CRID=”crid://example.com/abc123”

http://broadcaster.com/?CRID=”crid://broadcaster.com/abc123”&CRID=”crid://broadcaster.com/def456”

http://kaas.example.com/scripts/resolution?CRID=”crid://example.com/abc123”&Result=1

http://redirected.example.com/tva/lr?CRID=”crid://example.com/abc123”&SubmittedCRID=1

10.3.6.1 Further Requirements on a PDR HTTP Client.

The PDR shall at least implement the HTTP v1.0 specification [6] for issuing a GET request to the server. In addition to the requirements of HTTP v1.0, the PDR shall also send the HTTP v1.1 header of “host” and the HTTP v1.0 header “user-agent”.

The HTTP client in the PDR shall support at least the following MIME types:

- text/xml
- application/x-resolving-authority-record
- application/resolving-authority-record

E.g. the HTTP client will need to send an accept command with at least the following components:

Accept: text/xml; application/x-resolving-authority-record; application/resolving-authority-record

In order to permit the secure transfer of resolution requests from the PDR to the location resolution server and secure results from the location resolution server, the PDR and location resolution server may negotiate the secure HTTP protocol as given in RFC2660 [7].

10.3.7 TCP/IP Protocol for Response from Resolution Server

The response from the location resolution server will be based on the HTTP v1.0 specification [6].

The response from the server can be one of three possible types:

(1) The result of resolving the CRIDs sent to the server.

(2) An HTTP re-direct to allow the distribution of services amongst a number of machines.

(3) A standard Resolving Authority Record (RAR) to facilitate PDR caching, server load balancing, dynamic server administration and cross platform capability.

The use of a MIME type given by the "Content-Type" HTTP header shall be used to indicate which of the two possible server responses (type 1 or type 3) is being returned.

10.3.7.1 Case 1: Returning the Result of Resolving CRIDs

If the response from the location resolution server is the result of resolving the CRIDs requested by the PDR, one or more instances of XML Schema as defined in this specification or in TV Anytime specification SP003 [11] shall be returned.

The MIME type returned by the location resolution server shall be text/xml.

E.g. one of the response lines from the HTTP server will be:

Content-Type: text/xml

10.3.7.2 Case 2: Returning an HTTP Re-direct

Use of the HTTP redirect commands (HTTP error codes 300 to 399) may be used by the location resolution server to indicate that the PDR should disconnect and connect to a different location resolution server. The reason for providing this functionality is so that a location resolution provider can redirect their resolution requests based on the CRID being resolved rather than just the authority name (which can be re-directed during the DNS lookup phase of CRID resolution).

The "location" response header shall be used to indicate where the PDR should contact.

E.g.

Location: http://www.example.com/tvaresolve

When the PDR has been re-directed from its initial location resolution server, it should provide the HTTP v1.0 "`Referer`" header to the server to which it has been redirected.

10.3.7.3 Case 3: Returning a Resolving Authority Record.

In the case returning a resolving authority record, the MIME type shall be either application/x-resolving-authority-record or application/resolving-authority-record.

E.g. one of the response lines from the HTTP server will be:

Content-Type: application/resolving-authority-record

10.3.7.4 Encoding Server Response

It is permissible for the response from the server to be encoded, for example by compressing or encrypting the XML instance document. The "Content-Type" response header does not change, but rather the "Content-Encoding" gives the encoding that has been performed on the data. The exact form of encoding used is not specified in this specification and it is the responsibility of the HTTP client and server to negotiate commonly understood encoding systems.

E.g.

Content-Type: text/xml
Content-Encoding: gzip

11. References

[1] RFC2396 – Uniform Resource Identifiers (URI): Generic Syntax
T. Berners-Lee, R. Fielding, U.C. Irvine, L. Masinter
http://www.ietf.org/rfc/rfc2396.txt

[2] RFC1591 – Domain Name System Structure and Delegation
J. Postel
http://www.ietf.org/rfc/rfc1591.txt

[3] ANSI X3.4-1986 (R1997) – Information Systems - Coded Character Sets - 7-Bit American National Standard Code for Information Interchange (7-Bit ASCII) (reaffirmation of ANSI X3.4-1986 (R1992))

[4] RFC2068 – Hypertext Transfer Protocol -- HTTP/1.1
R. Fielding, UC Irvine, J. Gettys, J. Mogul, DEC, H. Frystyk, T. Berners-Lee, MIT/LCS
http://www.ietf.org/rfc/rfc2068.txt

[5] RFC2169 – A Trivial Convention for using HTTP in URN Resolution
R. Daniel
http://www.ietf.org/rfc/rfc2169.txt

[6] RFC1945 – Hypertext Transfer Protocol -- HTTP/1.0
T. Berners-Lee, R. Fielding & H. Frystyk.
http://www.ietf.org/rfc/rfc1945.txt

[7] RFC2660 – The Secure HyperText Transfer Protocol
E. Rescorla, A. Schiffman
http://www.ietf.org/rfc/rfc2660.txt

[8] RFC2782 – A DNS RR for Specifying the Location Of Services (DNS SRV)
A. Gulbrandsen, P.Vixie, L.Esibov.
http://www.ietf.org/rfc/rfc2782.txt

[9] HTML 4.01 Specification.
Ragget D., LeHors A., Jacobs L.
http://www.w3.org/TR/html4

[10] The WWW Common Gateway Interface version 1.1
Coar K.A.L et al.
http://www.golux.com/coar/cgi/draft-coar-cgi-v11-03-clean.html

NCSA Common Gateway Interface: http://hoohoo.ncsa.uiuc.edu/cgi/

[11] TV Anytime Specification SP003 v1.0
The TV Anytime Forum
http://www.tv-anytime.org

Appendix A : XML Schema for Content Referencing

A.1 Schema Definition

In this section, we define the normative content referencing schema. Instances of this schema are used during the location resolution process via bi-directional networks.

```
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema
 targetNamespace=http://www.tv-anytime.org/2001/04/ContentReferencing
 xmlns:metadata=http://www.tv-anytime.org/2001/04/metadata
 xmlns:xsd=http://www.w3.org/2000/10/XMLSchema
 xmlns:CR="http://www.tv-anytime.org/2001/04/ContentReferencing">
  <!-- Example import statement xsd:import
      namespace="http://www.tv-anytime.org/2001/04/metadata"
      schemaLocation="metadata.xsd"/> -->

      <xsd:element name="ContentReferencingTable"
      type="CR:ContentReferencingTableType">
            <xsd:annotation>
                  <xsd:documentation>A document conforming to the TV
                   Anytime content referencing specification
                  </xsd:documentation>
            </xsd:annotation>
      </xsd:element>

      <xsd:complexType name="ContentReferencingTableType">
            <xsd:sequence>
                  <xsd:element name="Result" type="CR:ResultType"
                   minOccurs="0" maxOccurs="unbounded"/>
            </xsd:sequence>
            <xsd:attribute name="version" type="xsd:float"
             use="required"/>
      </xsd:complexType>

      <xsd:simpleType name="AcquisitionDirectiveType">
            <xsd:restriction base="xsd:string">
                  <xsd:enumeration value="all"/>
                  <xsd:enumeration value="any"/>
            </xsd:restriction>
      </xsd:simpleType>

      <xsd:simpleType name="ResolutionStatusType">
            <xsd:restriction base="xsd:string">
                  <xsd:enumeration value="resolved"/>
                  <xsd:enumeration value="discard CRID"/>
                  <xsd:enumeration value="cannot yet resolve"/>
                  <xsd:enumeration value="unable to resolve"/>
            </xsd:restriction>
      </xsd:simpleType>

      <xsd:complexType name="ResultType">
            <xsd:choice>
                  <xsd:sequence>
                        <xsd:element name="CRIDResult"
                        type="CR:CRIDResultType" minOccurs="0"
```

```
                        maxOccurs="unbounded"/>
                  </xsd:sequence>
                  <xsd:sequence>
                        <xsd:element name="LocationsResult"
                         type="CR:LocationsResultType" minOccurs="0"
                         maxOccurs="unbounded"/>
                  </xsd:sequence>
            </xsd:choice>
            <xsd:attribute name="CRID" type="metadata:CRIDType"
             use="required"/>
            <xsd:attribute name="complete" type="xsd:boolean"
             use="required"/>
            <xsd:attribute name="acquire"
             type="CR:AcquisitionDirectiveType" use="required"/>
            <xsd:attribute name="status"
            type="CR:ResolutionStatusType" use="required"/>
            <xsd:attribute name="reresolveDate" type="xsd:dateTime"
                  use="optional"/>
      </xsd:complexType>

      <xsd:complexType name="CRIDResultType">
            <xsd:sequence>
                  <xsd:element name="Crid" type="metadata:CRIDType"
                        maxOccurs="unbounded"/>
            </xsd:sequence>
      </xsd:complexType>

      <xsd:complexType name="TimeAndURLType">
      <xsd:simpleContent>
            <xsd:extension base="xsd:uriReference">
            <xsd:attribute name="start" type="xsd:dateTime"
                  use="required"/>
            <xsd:attribute name="duration" type ="xsd:timeDuration"
                  use="optional"/>
            <xsd:attribute name="end" type="xsd:dateTime"
                  use="optional"/>
            </xsd:extension>
      </xsd:simpleContent>
      </xsd:complexType>

      <xsd:complexType name="LocationsResultType">
            <xsd:sequence>
            <xsd:choice>
                  <xsd:element name="Locator" type="xsd:uriReference"
                        maxOccurs="unbounded"/>
                  <xsd:element name="DecomposedLocator"
                   type="CR:TimeAndURLType" maxOccurs="unbounded"/>
                  </xsd:choice>
            </xsd:sequence>
      </xsd:complexType>
</xsd:schema>
```

Name	Definition
`ContentReferencingTable`	The top level element within which all content referencing results are instantiated.
`ContentReferencingTableType`	The syntax definition for the

	ContentReferencingTable element.
result	The content referencing information for each CRID being resolved are contained within this element.
version	The syntax version for this XML Schema. For instances conforming to the Schema defined in this specification, this field shall contain the value 1.0
ResultType	This type provides the container for all possible resolutions of a CRID.
CRID	The CRID that is being resolved
complete	True if the resolution of this CRID is complete. If false, the CRID may resolve into further CRIDs or locators in the future.
acquire	The grouping type for the list or CRIDs or Locators. This field is only meaningful when status equals resolved.
status	The status of resolving this CRID
reresolveDate	If status equals "cannot yet resolve", Or, complete equals false, this field contains the date and time when re-resolution should next be attempted. This date and time shall be unambiguous with respect to timezone.
CRIDResultType	When a CRID resolves into one or more CRIDs, an instance of CRIDResultType shall be used.
Crid	One of the "output" CRIDs.
TimeAndURLType	This extension to the uriReference type holds the URL pointing to the content and has attributes that contain the timing information required for acquisition.
start	If this is scheduled content, the date and time when the content will start. For on demand content, this is the time and date at which the content is first available.
duration	For scheduled content, this is the duration of the content. This element shall not be used for on demand content.
end	For scheduled content, this attribute shall not be used. For on demand content, this attribute contains the first time and date when the content is no longer available.
LocationsResultType	When a CRID resolves into one or more locators, an instance of LocationsResultType shall be used.
Locator	One of the "output" locators. This element is used when the locator contains the necessary timing information (such as start time, duration) as part of its syntax.
DecomposedLocator	One of the "output" locators. This element is used when the locator does not contain the necessary timing information (such as start time, duration) as

	part of its syntax.
AcquisitionDirectiveType	When a CRID resolves into a list of CRIDs or locators, the AcquisitionDirectiveType type describes what sort of group this list represents
all	All items in the list must be acquired. All items in the list combine to create the entire content assigned to the CRID being resolved.
any	One of the items from the list should be acquired. All items in the list are equivalent.
ResolutionStatusType	Type indicating the result of a resolution request.
resolved	The CRID has been successfully resolved
discard CRID	The CRID should be discarded
cannot yet resolve	The CRID cannot be resolved yet, and should be resolved again at a later date
unable to resolve	The CRID cannot be resolved.
CRIDResult	An element for instantiating a result that represents a CRID resolving into one or more CRIDs.
LocationsResult	An element for instantiating a result that represents a CRID resolving into one or more locations.

A.2 Example Instance Document

The following is an example instance document conforming to the XML Schema specified in this specification.

```
<?xml version="1.0" encoding="UTF-8"?>
<ContentReferencingTable
  xmlns="http://www.tv-anytime.org/2001/04/ContentReferencing"
 xmlns:xsi="http://www.w3.org/2000/10/XMLSchema-instance"
 xsi:schemaLocation="http://www.tv-anytime.org/2001/04/ContentReferencing
  local_copy_of_schema.xsd"
 version="1.0">

       <!-- Example of a CRID resolving to other CRIDs -->
       <Result CRID="crid://broadcaster.co.uk/akdsjdlkjdf" status="resolved"
                complete="true" acquire="all">
              <CRIDResult>
                     <Crid>CRID://example.com/greatstuff</Crid>
                     <Crid>CRID://nextcrid.com/lkjkj</Crid>
              </CRIDResult>
       </Result>

       <!-- Example of a CRID that is no longer valid -->
       <Result CRID="crid://isp.net/868457549845f" status="discard CRID"
        complete="true" acquire="all"/>

       <!-- Example of a CRID resolving to other CRIDs and is incomplete -->
       <Result CRID="crid://example.co.uk/wibble" status="resolved"
                complete="false" acquire="all">
              <CRIDResult>
                     <Crid>CRID://example.com/stuff</Crid>
                     <Crid>CRID://nextcrid.com/brodje</Crid>
              </CRIDResult>
       </Result>

       <!-- Example of a CRID resolving to locators -->
       <Result  CRID="crid://broadcaster.com/ajcnd" status="resolved" complete="true"
        acquire="all">
              <LocationsResult>
                     <Locator>
                            dvb://1.4ee2.3f4;4f5@2001-03-27T18:00:00.00+01:00
                     </Locator>
                     <DecomposedLocator start="2001-03-29T18:00:00.00">
                            ftp://myserver.example.com/directory12/hello.mp3
                     </DecomposedLocator>
                     <DecomposedLocator start="2001-03-29T18:00:00.00"
```

```
                                    end="2001-04-03T18:00:00.00">
                                ftp://myserver.example.com/directory12-backup/hello.mp3
                        </DecomposedLocator>
                </LocationsResult>
        </Result>
</ContentReferencingTable>
```

Appendix B : Example Dynamic Communication Behaviour between a PDR and a Remote Location Resolution Server

The following is an informative description of how the dynamic communication behaviour between a PDR and a remote Location Resolution Server may be implemented on the PDR.

If the response to location request is “resolve again after a given date and time”:

1) If the date and/or time is in the future:

   Wait until this date and time arrives. Wait for a further random amount of time before contacting the server.

2) Else If the date and time has already past:

   Wait for a random amount of time before contacting the server

   If the result from contacting the server again is that the re-resolve date and time is still in the past, double the range of the random delay from its current range and try again.

   When delay >= 1 day, the PDR can either:

   a) Try each day

   b) Continue doubling delay until delay >= one week, at which point the delay stays fixed at one week.

Else If the LR server is unavailable:

Wait a random amount of time and then try again. If server is still unavailable, double the range of random delay from its current range and try again.

When delay >= 1 day, the PDR can either:

a) Try each day

b) Continue doubling delay until delay >= one week, at which point the delay stays fixed at one week.

Else If response is “CRID is unknown” and server type = secondary:

Wait for a random amount of time before contacting the server.If the result from contacting the server again is that the crid is still unknown, double the range of the delay from its current range and try again.

When delay >= 1 day, the PDR can either:

a) Try each day

b) Continue doubling delay until delay >= one week, at which point the delay stays fixed at one week.

Else If response is “CRID is unknown” and server type = primary:

Either:

a) Assume CRID is invalid (as if LR server returns CRID = invalid)

Or:

b) Follow delay system above. After one day assume CRID is invalid.